

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
APPLICATION FOR LETTERS PATENT

LEAKAGE POWER ESTIMATION

INVENTORS:

Pedro Chaparro Monferrer

Grigorios Magklis

José González

Antonio González

ATTORNEY'S DOCKET NO. P23882PCT

LEAKAGE POWER ESTIMATION

BACKGROUND

[0001] The present disclosure generally relates to the field of electronics. More particularly, an embodiment of the invention relates to leakage power estimation in an integrated circuit (IC) device.

[0002] Power consumption, both dynamic and leakage, is one of the major concerns in IC design. In particular, sub-threshold leakage (or leakage power) may be growing with each successive design generation. For example, as supply voltage is lowered (e.g., to reduce dynamic power consumption), threshold voltage may also be lowered (e.g., to maintain low gate delay or high frequency). However, lowering the threshold voltage may affect leakage power nonlinearly.

[0003] In some implementations, leakage power may be assumed to have a constant value during run-time. However, leakage power may vary during run-time, for example, due to changes in temperature, supply voltage, or threshold voltage. Accordingly, power management techniques may be less accurate without knowledge of leakage power.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] The detailed description is provided with reference to the accompanying figures. In the figures, the left-most digit(s) of a reference number identifies the figure in which the reference number first appears. The use of the same reference numbers in different figures indicates similar or identical items.

[0005] Figs. 1, 5, and 6 illustrate block diagrams of computing systems in accordance with various embodiments of the invention.

[0006] Figs. 2A and 2B illustrate block diagrams of portions of leakage power estimation systems, according to various embodiments.

[0007] Fig. 3 illustrates a block diagram of a processor core, according to an embodiment.

[0008] Fig. 4 illustrates a flow diagram of a method, according to an embodiment.

DETAILED DESCRIPTION

[0009] In the following description, numerous specific details are set forth in order to provide a thorough understanding of various embodiments. However, various embodiments of the invention may be practiced without the specific details. In other instances, well-known methods, procedures, components, and circuits have not been described in detail so as not to obscure the particular embodiments of the invention. Various aspects of embodiments of the invention may be performed using various means, such as integrated semiconductor circuits (“hardware”), computer-readable instructions organized into one or more programs (“software”), or some combination of hardware and software. For the purposes of this disclosure reference to “logic” shall mean either hardware, software, or some combination thereof.

[0010] Some of the embodiments discussed herein may provide an efficient technique to estimate leakage power (e.g., static or a sub-threshold leakage power generated by one or more components of an IC device). In an embodiment, the leakage power consumption may be due to one or more variations such as variations in temperature and/or voltage (e.g., threshold and/or supply voltage). Furthermore, some of the embodiments discussed herein may be applied in various computing systems, such as the computing systems discussed with reference to Figs. 1, 5, and 6. More particularly, Fig. 1 illustrates a block diagram of a computing system 100, according to an embodiment. The system 100 may include one or more domains

102-1 through 102-M (collectively referred to herein as “domains 102” or “domain 102”). Each of the domains 102-1 through 102-M may include various components, but for clarity, sample components are only shown with reference to domains 102-1 and 102-2. Also, each domain 102 may correspond to a portion of a computing system (such as the components discussed with reference to Figs. 5 and 6, or more generally to one or more transistors of an IC device). In an embodiment, each of the domains 102 may include various circuitry (or logic) that is clocked by a clock signal that may be different than the clock signal used in other domains. In one embodiment, one or more of these clock signals may be mesosynchronous, or otherwise related (e.g., with a relationship that may or may not repeat itself over time).

[0011] As illustrated in Fig. 1, each domain may communicate data with other domains through one or more buffers 104. In an embodiment, the buffers 104 may be first-in, first-out (FIFO) buffers. Each domain may include a logic to estimate leakage power of one or more components within the corresponding domain (such as logics 106-1 and 106-2 shown with reference to domains 102-1 and 102-2, respectively, and generally referred to herein as “logic 106” or “logics 106”), one or more temperature sensors (such as sensor(s) 108-1 and 108-2 shown with reference to domains 102-1 and 102-2, respectively), a logic to control frequency and/or voltage levels and/or provide current threshold voltage and/or supply voltage values

(e.g., logics 110-1 and 110-2 shown with reference to domains 102-1 and 102-2, respectively), and a logic to manage power consumption of one or more components of the corresponding domain (such as logics 112-1 and 112-2 shown with reference to domains 102-1 and 102-2, respectively, and generally referred to herein as “logic 112” or “logics 112”). In an embodiment, the threshold voltage of a transistor may be adjusted by applying a current to the body (or substrate) of the transistor.

[0012] In various embodiments, the power management logic 112 may adjust power consumption of one or more components of a corresponding domain. For example, the logic 112 may utilize information such as the leakage power estimation value (e.g., provided by the corresponding logic 106), dynamic power estimation, and/or some other information (e.g., committed instructions per cycle, cache misses, etc.) to adjust supply voltage and/or threshold voltage of one or more components of the corresponding domain. Also, the logic 112 may adjust the frequency of a clock signal (e.g., a clock signal that is used within at least a portion of the corresponding domain). In an embodiment, the logic 112 may turn off one or more components such: one or more processor cores or portions of the processor cores (e.g., different pipelines, etc.) and/or data caches (e.g., including various levels of caches such as level 1 (L1), level 2 (L2), or other levels) or portions of data caches (e.g., different banks of caches).

[0013] Figs. 2A and 2B illustrate block diagrams of portions of leakage power estimation systems 200 and 250, according to various embodiments. In one embodiment, the systems 200 and 250 may be the same or similar to the logic 106 discussed with reference to Fig. 1. In an embodiment, the storage units discussed with reference to Figs. 2A and 2B may be the same or similar to memory components discussed with reference to Figs. 5 and/or 6.

[0014] As shown in Figs. 2A and 2B, the systems 200 and 250 may include a temperature scaling factor storage unit 202 (e.g., to store a plurality of temperature scaling factor values). The storage unit(s) 202 may receive sensed temperature values from the sensors 108 that correspond to one or more components such as those discussed with reference to Figs. 1, 5, and 6. The system 200 may also include a voltage scaling factor storage unit 204 (e.g., to store a plurality of voltage factor values) and a reference leakage storage unit 206 (e.g., to store a reference or base leakage power value). The base leakage value stored in the storage unit 206 may be determined at design time (e.g., through simulations or circuit measurements) or at test time. For example, the base leakage value may be determined at test time for designs where there is a relatively high variability (since the base value may be calculated independently for each chip and/or block to allow for adapting the estimations to the specifics of each circuit).

[0015] In an embodiment, the system 200 may also include a rounding logic 210 to round temperature values received from the sensors 108 (e.g., such that values sensed may be rounded to a nearest value stored in the storage unit 202). An interpolation logic 212 may interpolate the values output by the storage unit 202 to actual temperature measurement provided by the sensors 108. Similarly, the system 200 may include a voltage rounding logic 214 (e.g., to round current threshold and/or supply voltage values to a nearest value stored in the storage unit 204) and a voltage interpolation logic 218 (e.g., to interpolate the values output by the storage unit 204 to actual voltage values provided by the control logic 110). A multiplier 208 may multiply the determined temperature scaling factor (e.g., looked up from the storage unit 202 based on sensed temperature values from sensor(s) 108), the determined voltage scaling factor (e.g., looked up from the storage unit 204 based on current voltage values provided by logic 110), and the reference leakage value (from the storage unit 206). The multiplication value may then be utilized to manage power settings (e.g., by the power management logic 112) such as discussed with reference to Fig. 1.

[0016] Referring to Fig. 2B, the system 250 may include a reference leakage storage unit 252 that stores base leakage values for a corresponding set of voltages. Accordingly, in one embodiment, a single storage unit (252) may store values that correspond to a combination of values stored in the reference leakage storage unit

206 of Fig. 2A and corresponding values stored in the voltage scaling factor storage 204 of Fig. 2A. For example, a plurality of leakage power values may be indexed by a temperature factor (e.g., provided by the sensor(s) 108) and a voltage factor (e.g., corresponding to the threshold voltage value and/or supply voltage value provided by logic 110). Such an embodiment may allow a single look up (e.g., based on current threshold and/or supply voltage values from the logic 110) to provide a reference leakage value that may be scaled by the temperature scaling factor looked up from the storage unit 202 (e.g., based on sensed temperature value(s) provided by sensors 108) via a multiplier 254. Alternatively, the values stored in the storage units 202, 204, 206, and/or 252 may be combined into a single storage unit (not shown) to allow a single look up to provide a leakage value that corresponds to sensed temperature value(s) provided by sensors 108 and/or current threshold and/or supply voltage values from the logic 110. Also, the system 250 may include rounding and/or interpolation logic (e.g., that may be the same or similar to the logics 210, 212, 214, and/or 218) in accordance with some embodiments.

[0017] Fig. 3 illustrates a block diagram of a processor core 300, according to an embodiment. In one embodiment, the core 300 may represent various components that may be present in a processor or number of processors (such as those discussed with reference to Figs. 5 and 6). The processor core 300 may include one or more domains such as a second level cache domain 302, a frontend domain

304, and one or more backend domains 306. Components within each of the domains 302, 304, and 306 may be clocked by a different clock signal such as discussed with reference to Fig. 1. Moreover, each of the domains (e.g., 302, 304, and 306) may include more or less components than those shown in Fig. 3 in various embodiments.

[0018] The second level (L2) cache domain 302 may include an L2 cache 308 (e.g., to store data including instructions), the sensor(s) 108, and logics 106, 110, and 112. In one embodiment, the L2 cache 308 may be shared by multiple cores in a multi-core processor such as those discussed with reference to Figs. 5 and 6. Also, the L2 cache 308 may be off of the same die as the processor cores. Accordingly, in various embodiments of the invention, a processor may include the domains 304 and 306, and may or may not include the L2 cache 308.

[0019] As shown in Fig. 3, the frontend domain 304 may include one or more of the sensor(s) 108, logics 106, 110, and 112, a reorder buffer 318, a rename and steer unit 320, a instruction cache 322, a decode unit 324, a sequencer 326, and/or a branch prediction unit 328. In one embodiment, the frontend domain 304 may include other components such as an instruction fetch unit.

[0020] The backend domains 306 may include one or more of a first level (L1) cache domain 328 and one or more execution domains 330-1 through 330-N. The L1 cache domain 328 may include an L1 cache 332 (e.g., to store data including

instructions), the sensor(s) 108, and logics 106, 110, and 112. Furthermore, the execution domains 330-1 through 330-N may include one or more of an integer execution unit and/or a floating point execution unit. The execution domains 330-1 through 330-N may each comprise an issue queue (338-1 through 338-N, respectively), a register file (340-1 through 340-N, respectively), the sensor(s) 108, logics 106, 110, and 112, and/or an execution unit (346-1 through 346-N, respectively).

[0021] In one embodiment, each of the domains 302, 304, and 306 may include one or more first-in, first-out (FIFO) buffer(s) 348 to synchronize communication between the various clock domains (e.g., between the domains 302, 304, and/or 306).

[0022] Additionally, the processor core 300 (and, in an embodiment, such as the one shown in Fig. 3, the backend domains 306) may include an interconnection or bus 350 to facilitate communication between various components of the processor core 300. For example, after an instruction is successfully executed (e.g., by the execution domains 330-1 through 330-N), the instruction commit may be communicated to the ROB 318 (e.g., via the interconnection 350) to retire that instruction. Additionally, the domains within the backend (e.g., domains 328 and 330-1 through 330-N) may communicate via the interconnection 350. For example, communication among execution units (330-1 through 330-N) may occur for type

conversion instructions. Further operations of components of Figs. 1-3 will be discussed with reference to the method 400 of Fig. 4.

[0023] Furthermore, even though Fig. 3 illustrates that each of the domains 302, 304, and 306 may include the sensor(s) 108 and logics 106, 110, and 112, various domains may share the same the sensor(s) 108 and logics 106, 110, and 112. For example, a single set of the sensor(s) 108 and logics 106, 110, and 112 may be utilized for all domains of the processor core 300.

[0024] Fig. 4 illustrates a flow diagram of a method 400 to provide estimate leakage power, according to an embodiment. In one embodiment, the operations of the method 400 may be performed by one or more components, such as the components discussed with reference to Figs. 1-3 and 5-6.

[0025] Referring to Figs. 1-4, at an operation 402, the sensor(s) 108 may sense one or more temperature values corresponding to an IC device. The sensed temperature value(s) may be used to determine a temperature scaling factor (e.g., from the storage unit 202) at an operation 404. At operation 404, a voltage scaling factor may also be determined such as discussed with reference to Figs. 2A and 2B (e.g., from the storage units 204 and/or 252). At an operation 406, the determined scaling factors of operation 404 may then be used to scale a base leakage value (e.g., stored in the unit 206 and/or 252) such as discussed with reference to Figs. 2A and 2B. At an operation 408, a signal may be generated (e.g., by the multipliers 205 and

254) that correspond to an estimated leakage power of the IC device. As discussed with reference to Fig. 1, the estimated leakage power (408) may be used to adjust power consumption of one or more components of a computing system (e.g., systems discussed with reference to Figs. 1, 5, and/or 6).

[0026] In an embodiment, the following equation may be used to provide the estimate leakage power at operation 408:

$$P(V_{dd}, V_{th}, T) = P_0 \cdot \frac{V_{dd}}{V_{dd0}} \cdot e^{\beta(V_{dd} - V_{dd0})} \cdot e^{\gamma \cdot (-|V_{th}| + |V_{th0}|)} \cdot e^{\delta(T - T_0)}$$

[0027] In the above formula, P corresponds to the estimate leakage power value, P_0 corresponds to the base leakage power value (e.g., that may be stored in units 206 and/or 252), V_{dd} corresponds to supply voltage (that may be provided by the logic 110), V_{th} corresponds to threshold voltage (that may be provided by the logic 110), V_{dd0} corresponds to V_{dd} at which base leakage was measured, V_{th0} corresponds to V_{th} at which base leakage was measured, T corresponds to current temperature value(s) sensed by the sensor(s) 108, T_0 corresponds to the temperature at which base leakage was measured, δ , β and γ are circuit dependent constants set by the designer. In various embodiments, values corresponding to term $T(T) = e^{\delta \cdot (T - T_0)}$ may be stored in the storage unit 202 and values corresponding to the term $V(V_{dd}, V_{th}) = \frac{V_{dd}}{V_{dd0}} \cdot e^{\beta(V_{dd} - V_{dd0})} \cdot e^{\gamma \cdot (-|V_{th}| + |V_{th0}|)}$ may be stored in the storage units 204

(or 252). Hence, a multiplier (208, 254) may be used to multiply the terms $T(T)$ and $V(V_{dd} V_{th})$ to provide value of P .

[0028] Moreover, in one embodiment, dynamic calibration of an IC component may be performed in idle mode (e.g., where there is no dynamic power consumption). In such situation, the temperature increase (over a controlled ambient temperature) in each portion (e.g., blocks) of the IC component may be dependant upon the leakage power. The thermal sensors 108 that may be placed in the blocks can report the stable temperature (e.g., after a relatively long period of time). With the temperature map, a tool (such as a computing device that is external to the IC component) may derive the power map that is causing the scenario, e.g., via reverse-engineering. The leakage values may then be computed based on the static temperatures of the portions (since other constants may be known, such as supply voltage, threshold voltage, and ambient temperature). Once the power map is computed it may be stored in the reference leakage storage 206. In an embodiment, a special dedicated microcode may be used to communicate between the IC component being calibrated and test equipment (e.g., to report the temperature readings and to perform the base leakage update).

[0029] Fig. 5 illustrates a block diagram of a computing system 500 in accordance with an embodiment of the invention. The computing system 500 may include one or more central processing unit(s) (CPUs) 502 or processors that

communicate via an interconnection network (or bus) 504. The processors 502 may be any type of a processor such as a general purpose processor, a network processor (that processes data communicated over a computer network 503), or other types of a processor (including a reduced instruction set computer (RISC) processor or a complex instruction set computer (CISC)). Moreover, the processors 502 may have a single or multiple core design. The processors 502 with a multiple core design may integrate different types of processor cores on the same integrated circuit (IC) die. Also, the processors 502 with a multiple core design may be implemented as symmetrical or asymmetrical multiprocessors. In an embodiment, one or more of the processors 502 may utilize the embodiments discussed with reference to Figs. 1-4. For example, one or more of the processors 502 may include one or more processor cores (300). Also, the operations discussed with reference to Figs. 1-4 may be performed by one or more components of the system 500.

[0030] A chipset 506 may also communicate with the interconnection network 504. The chipset 506 may include a memory control hub (MCH) 508. The MCH 508 may include a memory controller 510 that communicates with a memory 512. The memory 512 may store data and sequences of instructions that are executed by the CPU 502, or any other device included in the computing system 500. In one embodiment of the invention, the memory 512 may include one or more volatile storage (or memory) devices such as random access memory (RAM), dynamic RAM

(DRAM), synchronous DRAM (SDRAM), static RAM (SRAM), or the like. Nonvolatile memory may also be utilized such as a hard disk. Additional devices may communicate via the interconnection network 504, such as multiple CPUs and/or multiple system memories.

[0031] The MCH 508 may also include a graphics interface 514 that communicates with a graphics accelerator 516. In one embodiment of the invention, the graphics interface 514 may communicate with the graphics accelerator 516 via an accelerated graphics port (AGP). In an embodiment of the invention, a display (such as a flat panel display) may communicate with the graphics interface 514 through, for example, a signal converter that translates a digital representation of an image stored in a storage device such as video memory or system memory into display signals that are interpreted and displayed by the display. The display signals produced by the display device may pass through various control devices before being interpreted by and subsequently displayed on the display.

[0032] A hub interface 518 may allow the MCH 508 to communicate with an input/output control hub (ICH) 520. The ICH 520 may provide an interface to I/O devices that communicate with components of the computing system 500. The ICH 520 may communicate with a bus 522 through a peripheral bridge (or controller) 524, such as a peripheral component interconnect (PCI) bridge, a universal serial bus (USB) controller, or the like. The bridge 524 may provide a data path between the

CPU 502 and peripheral devices. Other types of topologies may be utilized. Also, multiple buses may communicate with the ICH 520, e.g., through multiple bridges or controllers. Moreover, other peripherals in communication with the ICH 520 may include, in various embodiments of the invention, integrated drive electronics (IDE) or small computer system interface (SCSI) hard drive(s), USB port(s), a keyboard, a mouse, parallel port(s), serial port(s), floppy disk drive(s), digital output support (e.g., digital video interface (DVI)), or the like.

[0033] The bus 522 may communicate with an audio device 526, one or more disk drive(s) 528, and a network interface device 530 (which communicates with the computer network 503). Other devices may be in communication with the bus 522. Also, various components (such as the network interface device 530) may be in communication with the MCH 508 in some embodiments of the invention. In addition, the processor 502 and the MCH 508 may be combined to form a single chip. Furthermore, the graphics accelerator 516 may be included within the MCH 508 in other embodiments of the invention.

[0034] Furthermore, the computing system 500 may include volatile and/or nonvolatile memory (or storage). For example, nonvolatile memory may include one or more of the following: read-only memory (ROM), programmable ROM (PROM), erasable PROM (EPROM), electrically EPROM (EEPROM), a disk drive (e.g., 528), a floppy disk, a compact disk ROM (CD-ROM), a digital versatile disk

(DVD), flash memory, a magneto-optical disk, or other types of nonvolatile machine-readable media capable of storing electronic instructions and/or data.

[0035] Fig. 6 illustrates a computing system 600 that is arranged in a point-to-point (PtP) configuration, according to an embodiment of the invention. In particular, Fig. 6 shows a system where processors, memory, and input/output devices are interconnected by a number of point-to-point interfaces. The operations discussed with reference to Figs. 1-5 may be performed by one or more components of the system 600.

[0036] As illustrated in Fig. 6, the system 600 may include several processors, of which only two, processors 602 and 604 are shown for clarity. The processors 602 and 604 may each include a local memory controller hub (MCH) 606 and 608 to allow communication with memories 610 and 612. The memories 610 and/or 612 may store various data such as those discussed with reference to the memory 512.

[0037] The processors 602 and 604 may be any type of a processor such as those discussed with reference to the processors 502 of Fig. 5. The processors 602 and 604 may exchange data via a point-to-point (PtP) interface 614 using PtP interface circuits 616 and 618, respectively. The processors 602 and 604 may each exchange data with a chipset 620 via individual PtP interfaces 622 and 624 using point to point interface circuits 626, 628, 630, and 632. The chipset 620 may also

exchange data with a high-performance graphics circuit 634 via a high-performance graphics interface 636, using a PtP interface circuit 637.

[0038] At least one embodiment of the invention may be provided within the processors 602 and 604. For example, one or more of the domains 102 discussed with reference to Fig. 1 and/or processor core(s) 300 may be located within the processors 602 and 604. Other embodiments of the invention, however, may exist in other circuits, logic units, or devices within the system 600 of Fig. 6. Furthermore, other embodiments of the invention may be distributed throughout several circuits, logic units, or devices illustrated in Fig. 6.

[0039] The chipset 620 may communicate with a bus 640 using a PtP interface circuit 641. The bus 640 may have one or more devices that communicate with it, such as a bus bridge 642 and I/O devices 643. Via a bus 644, the bus bridge 643 may be in communication with other devices such as a keyboard/mouse 645, communication devices 646 (such as modems, network interface devices, etc. that may be in communication with the computer network 503), audio I/O device, and/or a data storage device 648. The data storage device 648 may store code 649 that may be executed by the processors 602 and/or 604.

[0040] In various embodiments of the invention, the operations discussed herein, e.g., with reference to Figs. 1-6, may be implemented by hardware (e.g., circuitry), software, firmware, microcode, or combinations thereof, which may be

provided as a computer program product, e.g., including a machine-readable or computer-readable medium having stored thereon instructions (or software procedures) used to program a computer to perform a process discussed herein. Also, the term “logic” may include, by way of example, software, hardware, or combinations of software and hardware. The machine-readable medium may include a storage device such as those discussed with respect to Figs. 1-6. Additionally, such computer-readable media may be downloaded as a computer program product, wherein the program may be transferred from a remote computer (e.g., a server) to a requesting computer (e.g., a client) by way of data signals embodied in a carrier wave or other propagation medium via a communication link (e.g., a bus, a modem, or a network connection). Accordingly, herein, a carrier wave shall be regarded as comprising a machine-readable medium.

[0041] Reference in the specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment may be included in at least an implementation. The appearances of the phrase “in one embodiment” in various places in the specification may or may not be all referring to the same embodiment.

[0042] Also, in the description and claims, the terms “coupled” and “connected,” along with their derivatives, may be used. In some embodiments of the invention, “connected” may be used to indicate that two or more elements are in

direct physical or electrical contact with each other. "Coupled" may mean that two or more elements are in direct physical or electrical contact. However, "coupled" may also mean that two or more elements may not be in direct contact with each other, but may still cooperate or interact with each other.

[0043] Thus, although embodiments of the invention have been described in language specific to structural features and/or methodological acts, it is to be understood that claimed subject matter may not be limited to the specific features or acts described. Rather, the specific features and acts are disclosed as sample forms of implementing the claimed subject matter.